

MIT OpenCourseWare
<http://ocw.mit.edu>

14.123 Microeconomic Theory III
Spring 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.

MIT 14.123 (2009) by Peter Eso

Lecture 8: Equilibrium Refinements

1. Sequential Equilibrium
2. Trembling-Hand Perfect Equilibrium
3. Stability

Read: FT Chapter 8.




Previously in 14.123...

- We learned that (Correlated) Rationalizability is equivalent to Iterated Deletion of Strictly Dominated Strategies.
- Investigated subgame perfection in various classes of games.
 - In finite games with perfect information, use backward induction to find unique SPE. (We all knew that.)
 - In finite or infinite (w/ continuity at ∞) games with “almost perfect” information, use single-deviation principle.
 - In finite or infinite games with perfect information (e.g., alt. offer bargaining): iterated conditional dominance.
- Today: Refinements of subgame perfection in extensive-form games with imperfect or incomplete information.



Why Refine SPE?



Example 1.

- Single subgame, hence any Nash equilibrium is SPE.

- But what explains P2's choice of L ?

- If P2 is called upon to play, P1 played either U or D . In either case, R is best response.


At P2's only information set, no matter what he believes about P1's prior move, L is not rational given those beliefs.

*Aside: What if we changed $u((D,R))$ to $(3,1)$ and reduced the game?

Sequential Rationality

- A player is called sequentially rational (at a history) if s/he plays a best reply to a belief conditional on being at that history.
- DEF: An assessment is (σ, μ) where σ is a strategy profile and μ is a belief system; $\mu(h) \in \Delta(h)$ for all information sets h .
- DEF: An assessment (σ, μ) is sequentially rational if at each h_i , σ_i is a best reply to σ_{-i} given $\mu(h)$:
 $(R) \quad \forall i, \forall h_i, \forall s_i: E_{\mu(h_i)}[u_i(\sigma_i, \sigma_{-i})|h_i] \geq E_{\mu(h_i)}[u_i(s_i, \sigma_{-i})|h_i].$
- DEF: Belief-system μ is computed from strategy-profile σ (on the path of play) using Bayes' rule if
 $(B) \quad \forall h \text{ with } \sigma(h) > 0, \forall x \in h: \mu(x|h) = \Pr(x|\sigma) / \sum_{y \in h} \Pr(y|\sigma).$
- Is it enough to require Bayes rule and sequential rationality?

An Entry Deterrence Game



Example 2.

P1 playing *Out* and P2 playing *F* while holding beliefs $\mu_2 = 0$ is SPE.

Implausible for “technical” reason: P1 does not observe Nature’s move, so how could his action convey information (i.e, flip P2’s prior) about it?

*Aside: Rewrite the tree so that Nature moves last.



Sequential Equilibrium

Strengthen the Bayes-rule requirement (B) on page 4 as follows.

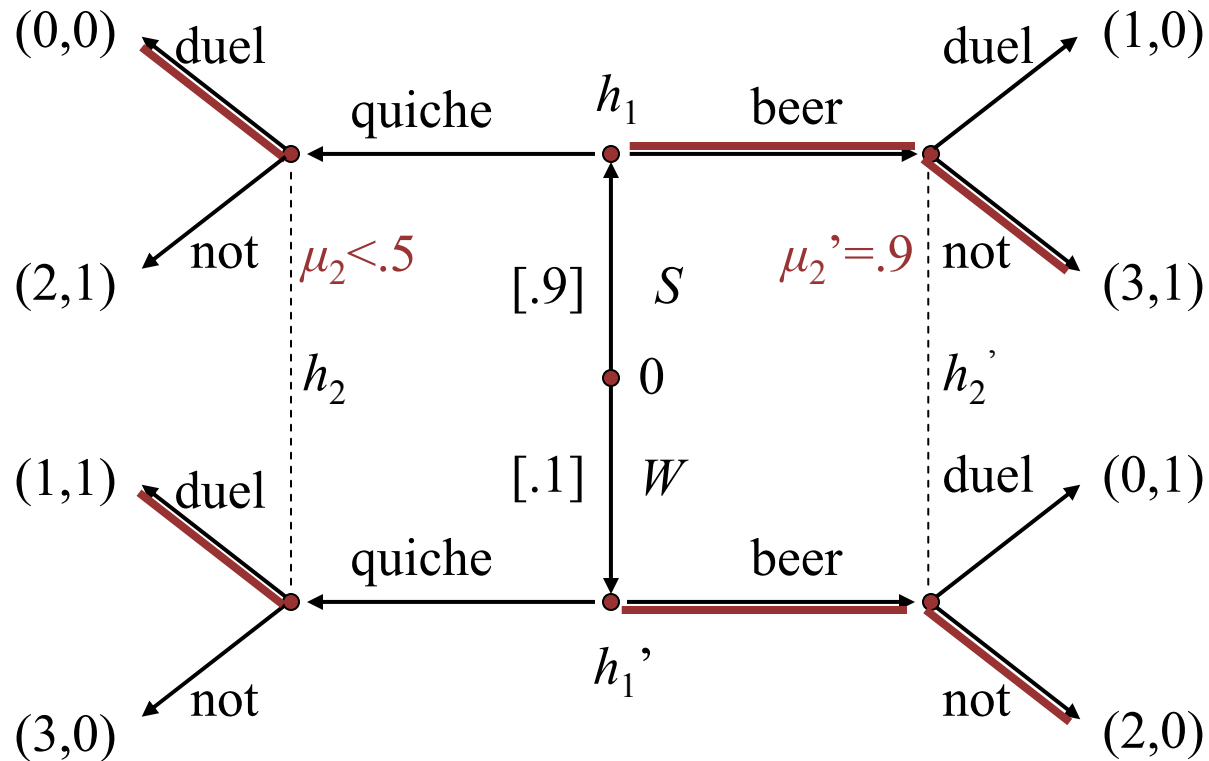
- DEF: Assessment (σ, μ) is consistent if there exists a sequence $(\sigma^m, \mu^m) \rightarrow (\sigma, \mu)$ such that σ^m is completely mixed and μ^m is computed from σ^m using Bayes' rule:

$$(C) \quad \forall h, x \in h: \mu^m(x|h) = \Pr(x|\sigma^m) / \sum_{y \in h} \Pr(y|\sigma^m).$$

- DEF: An assessment (σ, μ) is sequential equilibrium if it is sequentially rational (R), and consistent (C).
- THM: Sequential equilibrium exists in finite games and satisfies subgame perfection. It is upper-hemicontinuous in payoffs. The definition and theorem are due to Kreps and Wilson (1982).
- In Example 2, (Out, F) is not sequential equilibrium. (Why?)

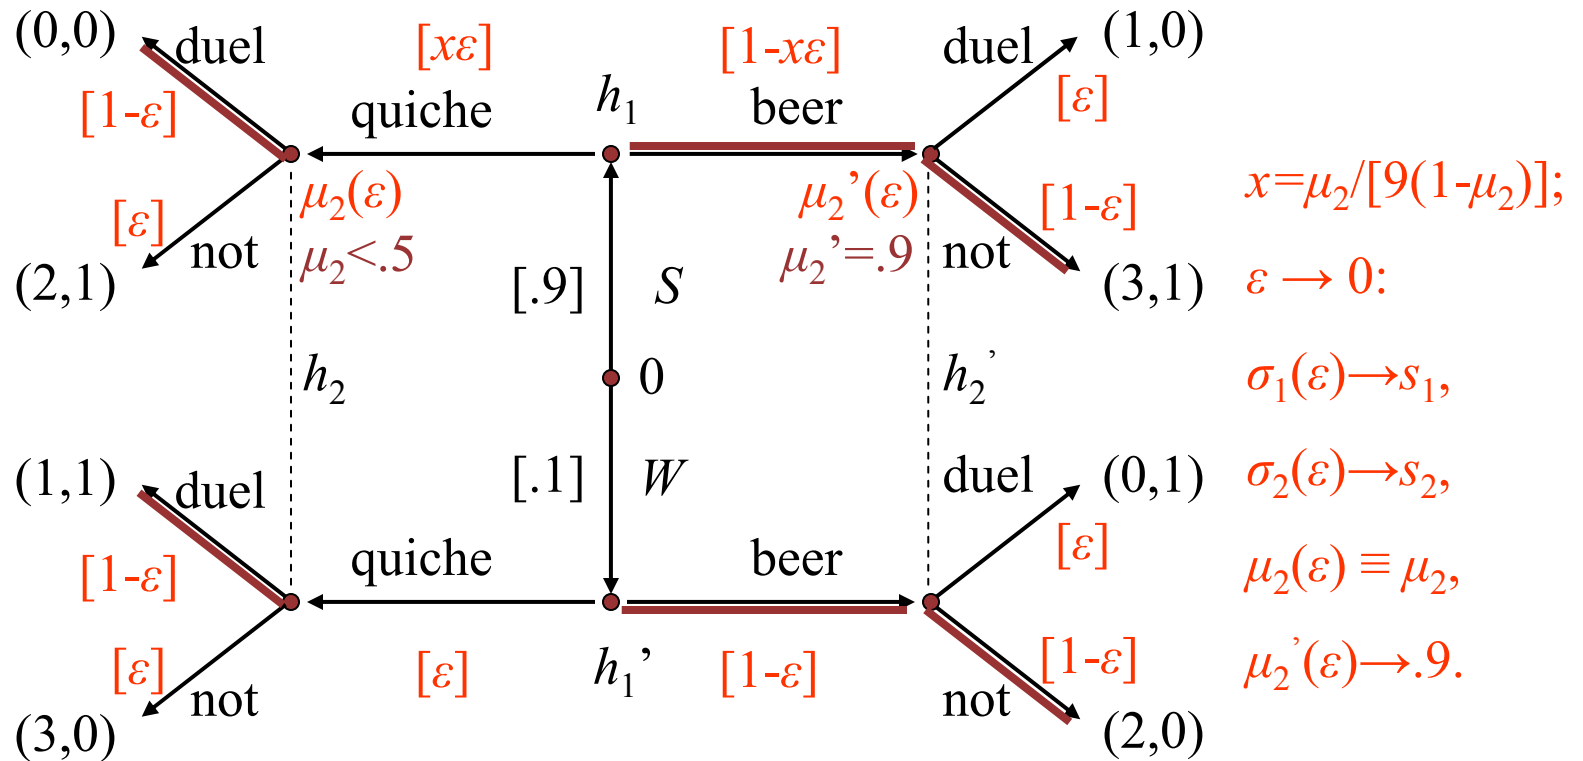


Sequential Eqm in Beer-Quiche




$$s_1 = (h_1, h_1' \mapsto beer), s_2 = (h_2 \mapsto duel, h_2' \mapsto not).$$

Sequential Eqm in Beer-Quiche



$$s_1 = (h_1, h_1' \mapsto \text{beer}), s_2 = (h_2 \mapsto \text{duel}, h_2' \mapsto \text{not}).$$

More SE with the Same Outcome




$$s_1 = (h_1, h_1' \mapsto \text{beer}), s_2 = (h_2 \mapsto \alpha \geq 1/2, h_2' \mapsto \text{not}).$$

“Equilibrium component”



Pooling on Quiche



$$s_1 = (h_1, h_1' \mapsto \text{quiche}), s_2 = (h_2 \mapsto \text{not}, h_2' \mapsto \text{duel}).$$

Trembling-Hand Perfection

- A stronger concept based on small mistakes (trembles) in the normal form, due to Selten (1975), before sequential equilibrium.
- DEF: σ is a trembling-hand perfect equilibrium if there exist a sequence $\sigma^m \rightarrow \sigma$ such that each σ^m is totally mixed, and for all i and $s_i \in S_i$: $u_i(\sigma_i, \sigma^m_{-i}) \geq u_i(s_i, \sigma^m_{-i})$.
- Alternatively: Let an ε -constrained equilibrium be a totally mixed σ^ε such that σ^ε_i is a best response to σ^ε_{-i} subject to $\sigma^\varepsilon_i(s_i) \geq \varepsilon(s_i)$ for some $\varepsilon(s_i) \in (0, \varepsilon)$, where $s_i \in S_i$, $i \in N$. A trembling-hand perfect equilibrium is any limit of ε -constrained equilibria as $\varepsilon \rightarrow 0$.
- Note that in both cases, it is enough to find just one converging sequence ($\sigma^m \rightarrow \sigma$ and $\sigma^\varepsilon \rightarrow \sigma$, resp.) with the required properties.

Trembling-Hand Perfection

- THM: In a finite game, trembling-hand equilibrium exists.
- $\Delta^\varepsilon(S_i) := \{ \sigma_i \in \Delta(S_i) \mid \sigma_i(s_i) \geq \varepsilon \text{ for all } s_i \in S_i \}$: compact, convex.
 $\Psi_i^\varepsilon(\sigma_{-i}) := \{ \sigma_i \in \Delta^\varepsilon(S_i) \mid u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}), \forall \sigma'_i \in \Delta^\varepsilon(S_i) \}$:
restricted best response, upper-hemicontinuous.
 $(\sigma_1, \dots, \sigma_n) \mapsto (\Psi_1^\varepsilon(\sigma_{-1}), \dots, \Psi_n^\varepsilon(\sigma_{-n}))$ has fixed point (Kakutani), in
which non-optimal strategies get exactly ε weight.
Any fixed point, σ^ε , is an ε -constrained equilibrium; set of fixed
points varies continuously with ε . Limit (in weak convergence)
is trembling-hand perfect equilibrium. ■

THP in Examples 1-2

	<i>L</i>	<i>R</i>
<i>U</i>	0,0	3,1
<i>D</i>	0,0	1,3
<i>X</i>	2,2	2,2

Example 1: (*X,L*) is SPE in extensive form, not sequential equilibrium.

L is weakly dominated, hence played with prob ε in any ε -constrained eqm; σ_2^ε cannot converge to *L* as $\varepsilon \rightarrow 0$, (*X,L*) is not THP.

Example 2: (*Out,F*) is SPE, not sequential equilibrium.

Normal form: P1 chooses rows, P2 columns, Nature matrices.

[.99]

	<i>F</i>	<i>A</i>
<i>In</i>	0,0	3,1
<i>Out</i>	2,2	2,2

[.01]

	<i>F</i>	<i>A</i>
<i>In</i>	0,0	1,1
<i>Out</i>	2,2	2,2

Expected payoffs \approx left matrix entries; only (*In,A*) is THP.


THP and Sequential Equilibrium

- THM: In any finite game where each player has only one move in the extensive form, all trembling-hand perfect equilibria are sequential. If the payoffs are generic (essentially: no indifference for any player), then the two concepts are equivalent.
 - Trembling-hand perfect strategies are limits of totally mixed strategies that define unique beliefs at every info set. The limits of these beliefs are consistent with the eqm strategies. Sequential rationality follows because player i , deciding at his only information set, gives a best response to σ_{-i} . ■
- In “agent normal form” (players represented by distinct agents at each info set, maximizing the same utility), for generic games, trembling-hand perfection is equivalent to sequential equilibrium.



Why Agent-Normal Form?

- Example 3. Unique SPE, unique sequential eqm: (UU',L) .



	L	R
UU'	3,1	1,0
UD'	0,-5	1,0
DD'	2,2	2,2

- However, (DD',R) is Nash, and trembling-hand perfect.
 - P1 plays UU' with ε^2 , UD' with ε , and DD' with $1-\varepsilon-\varepsilon^2$.
 P1's prob of playing D' conditional on U is $\varepsilon/(\varepsilon+\varepsilon^2) \approx 1$,
 hence P2's best response is to put max weight on R . ■
- Agent-normal form rules out such “correlated trembles”, makes THP as strong as SPE and sequential equilibrium.


Why Generic Payoffs?

- Example 4. (U,L) is the only trembling-hand perfect equilibrium; both (U,L) and weakly-dominated (D,R) are sequential equilibria.
- THP stronger than sequential eqm because it requires that σ^ε is an ε -equilibrium.
- If the payoffs are perturbed (so that each player has strict prefs over his actions given any action of the other), the equivalence of trembling-hand perfect and sequential equilibria is restored.
- Note: The set of trembling-hand perfect equilibria is not upper-hemicontinuous in payoffs (unlike set of sequential equilibria).
 - Replace $(0,0)$ with $(\delta,\delta) > 0$ as the payoffs of (D,R) in Ex.4. (D,R) is trembling-hand perfect $\forall \delta > 0$ but not for $\delta = 0$. ■

	<i>L</i>	<i>R</i>
<i>U</i>	1,1	0,0
<i>D</i>	0,0	0,0

Perfection is Imperfect

- Example 5. (X,L) is sequential eqm and trembling-hand perfect. P2 has 50-50% beliefs, plays L .
- ε -trembles: P1 plays U and D each with ε , P2 plays R with ε .
- But U strictly dominates D , why tremble with the same ε ?
- *Another issue: Add irrelevant move NX at initial move of P1 leading to P1 deciding U or D . (X,L) is not sequential / THP; unique SPE is $(NX-U, R)$.



		$[1-\varepsilon]$ L	R $[\varepsilon]$
$[\varepsilon]$ U		1,0	3,1
$[\varepsilon]$ D		0,2	0,0
$[1-2\varepsilon]$ X		2,2	2,2

‘Pooling on Quiche’ is THP

- Beer-quiche: Nature picks row; type of P1 matrix; P2 columns.

	duel	<u>not</u>	<u>duel</u>	not	prior
<i>S</i>	0,0	2,1	1,0	3,1	.9
<i>W</i>	1,1	3,0	0,1	2,0	.1
	<u>quiche</u>		beer		

- Trembles used to validate consistency in sequential equilibrium (page 10) prove that ‘Pool on quiche’ is trembling-had perfect.
- But the only agent of P1 that can improve his equilibrium payoff by playing ‘beer’ is type *S*; type *W* already gets his best outcome.
- Forward induction: P2 tries to find out what P1 wants to “say” by going off-equilibrium – P2 does not think it is a mistake!

Strategic Stability

- Ex.5 and ‘pooling on quiche’ in beer-quiche suggest we should require robustness to all trembles. Kohlberg & Mertens (1986).
- DEF: σ^* is strategically stable if, for any sequence of $\{\varepsilon^m(s_i) > 0 : s_i \in S_i, i \in N\}$ with $\varepsilon^m \rightarrow 0$ as $m \rightarrow \infty$, there exists a sequence $\sigma^m \rightarrow \sigma^*$ such that σ^m_i is a best response to σ^m_{-i} subject to $\sigma^m_i(s_i) \geq \varepsilon^m(s_i)$.
- Stability is stronger than trembling-hand perfection because we require convergence to σ^* for any (not some) collection $\varepsilon^m(s_i) \rightarrow 0$.
- Stability checks if the equilibrium is robust to any trembles.
Is this really what forward induction means? FT pp. 464-466.
- In ‘beer-quiche’ game, assign $\varepsilon_S \gg \varepsilon_W$ to show that ‘pooling on quiche’ is not strategically stable.

Non-Existence

- Two equilibrium components: (U,L) and $\sigma^p = (X, pL + (1-p)R)$ with $p \leq 1/2$.
- No equilibrium is stable in this example.
- (U,L) is not even trembling-hand perfect because U is weakly dominated.

	L	R
X	2,0	2,0
U	3,1	0,0
D	3,1	1,2

Suppose $\varepsilon_U > \varepsilon_D$ and $\sigma^m \rightarrow \sigma^p$. If D is a best response to σ^m_2 , then $\sigma^m_2(L) \geq 1/2$ and so $p = 1/2$ in the limit. If D is not P1's best response, then D is played with prob ε_D (while U is played more often), and so P2's best response to σ^m_1 is L , contradicting $\sigma^m \rightarrow \sigma^p, p \leq 1/2$.

Thus, the only candidate for a stable equilibrium is σ^p with $p = 1/2$.

Now consider $\varepsilon_D > \varepsilon_U$, $\sigma^m \rightarrow \sigma^p$. P2's unique best response is R unless P1's best response is U , which is impossible with $p = 1/2$. ■

Strategic Stability, Take 2

- DEF: A set of Nash equilibria is stable if it is minimal w.r.t. the property that for any sequence of positive weights $\varepsilon^m(s_i) \rightarrow 0$, there exists a sequence σ^m such that σ^m_i is a best response to σ^m_{-i} subject to $\sigma^m_i(s_i) \geq \varepsilon^m(s_i)$, and σ^m converges to a point in the set.
- In the example, the set $\{(X, \frac{1}{2}L + \frac{1}{2}R), (X, R)\}$ is a stable set.
 - For some sequences the limit is σ^p with $p = \frac{1}{2}$, for others it is σ^p with $p = 1$. Compare with the nonexistence proof.■
- THM (Kohlberg & Mertens): Stable sets exist in finite games. Stable sets coincide for all extensive-form games with the same normal form. No σ^* in a stable set assigns positive probability to a weakly dominated strategy. Every stable set of G contains a stable set of G' obtained from G by iterated strict dominance.

*Refine Stability?

	<i>A</i>	<i>HR</i>	<i>CFO</i>	<i>A</i>	π
<i>Quant</i>	2,2	0,0	5,5	3,3	.5
<i>Poet</i>	2,2	1,5	0,0	3,3	.5
	<i>no MBA</i>	<i>MBA</i>			

- Nature picks row, P1 matrix, P2 column. Three pure equilibria, one is pooling on “*no MBA*” supported by *HR*, $\mu(\textit{Poet}|\textit{MBA}) > 3/5$.
- This equilibrium is stable (proof is HW). Is it “reasonable”?
 - P1 could deviate to *MBA* and say: “Believe me, I am a *Quant*. If you do, and play *CFO* (best reply), then no type but a *Quant* gains from this deviation. There is no similar speech for a *Poet*.”
- This is where we stop – but the literature goes on.